arXiv:1603.00937v1 [physics.comp-ph] 3 Mar 2016

Graph-based linear scaling electronic structure theory

Anders M. N. Niklasson,[1] Susan M. Mniszewski,[2] Christian F. A. Negre,[1] Marc J. Cawkwell,[1] Pieter J. Swart,[1] Jamal Mohd-Yusof,[2] Timothy C. Germann,[1] Michael E. Wall,[2] Nicolas Bock,[1] Emanuel H. Rubensson,[3] and Hristo Djidjev[4]

[1] *Theoretical Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545*

[2] *Computer, Computational, and Statistical Sciences Division, Los Alamos National Laboratory, Los Alamos, New Mexico 8754*

[3] *Division of Scientific Computing, Department of Information Technology, Uppsala University, Box 337, SE-751 05 Uppsala, Sweden*

[4] *Computer, Computational, and Statistical Sciences Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545*

(Dated: March 9, 2016)

Abstract

We show how graph theory can be combined with quantum theory to calculate the electronic structure of large complex systems. The graph formalism is general and applicable to a broad range of electronic structure methods and materials, including challenging systems such as biomolecules. The methodology combines well-controlled accuracy, low computational cost, and natural low-communication parallelism. This combination addresses substantial shortcomings of linear scaling electronic structure theory, in particular with respect to quantum-based molecular dynamics simulations.

I. INTRODUCTION

The importance of electronic structure theory in materials science, chemistry, and molecular biology, relies on the development of theoretical methods that provide sufficient accuracy at a reasonable computational cost. Currently, the field is dominated by Kohn-Sham density functional theory[1–4], which often combines good theoretical fidelity with a modest computational workload that is constrained mainly by the diagonalization of the Kohn-Sham Hamiltonian – an operation that scales cubically with the system size. However, for systems beyond a few hundred atoms, the diagonalization becomes prohibitively expensive. This bottleneck was removed with the development of linear scaling electronic structure theory[5,6], which allows calculations of systems with millions of atoms[7,8]. Unfortunately, the immense promise of linear scaling electronic structure theory has never been fully realized because of some significant shortcomings, in particular, *a*) the accuracy is reduced to a level that is often difficult, if not impossible, to control; *b*) the computational pre-factor is high and the linear scaling benefit occurs only for very large systems that in practice often are beyond acceptable time limits or available computer resources; and *c*) the parallel performance is generally challenged by a significant overhead and the wall-clock time remains high even with massive parallelism. In quantum-based molecular dynamics (QMD) simulations[9] all these problems coalesce and we are constrained either to small system sizes or short simulation times.

In this paper we propose to overcome these shortcomings by introducing a formalism based on graph theory[10,11] that allows practical and easily parallelizable electronic structure calculations of large complex systems with well-controlled accuracy. The graph-based electronic structure theory combines the natural parallelism of a divide and conquer approach[12–17] with the automatically adaptive and tunable accuracy of a thresholded sparse matrix algebra[18–31], which can be combined with fast, low pre-factor, recursive Fermi operator expansion methods[32–41] and be applied to modern formulations of Born-Oppenheimer molecular dynamics (MD)[42–50].

II. RESULTS

A. Graph-based electronic structure theory

Our graph-based electronic structure theory relies on the equivalence between the calculation of thresholded sparse matrix polynomials and a graph partitioning approach. Let $P(X)$ be a Mth-order polynomial of a $N \times N$ symmetric square matrix X that is given as a linear combination of some basis polynomials $T^{(n)}(X)$,

$$P(X) = \sum_{n=0}^{M} c_n T^{(n)}(X). \tag{1}$$

We define an approximation $P_\tau(X)$ of $P(X)$ using a *globally* thresholded sparse matrix algebra, where matrix elements with a magnitude below a numerical threshold τ *in all terms*, $T^{(n)}(X)$, are ignored. The pattern of the remaining matrix entries, which at any point of the expansion have been (or are expected to be) greater than τ, can be described by a *data dependency graph* $\mathcal{S}_\tau$ that represents all possible data dependencies between the matrix elements in the polynomial expansion. Formally, we define the graph S_τ with a vertex for each row of X and an edge (i, j) between vertices i and j if

$$\{T^{(n)}(X)\}_{i,j} \geq \tau \text{ for any } n \leq M. \tag{2}$$

For a matrix A, we denote by $\lfloor A \rfloor_{\mathcal{S}_\tau}$ the thresholded version of A, where

$$\{\lfloor A \rfloor_{\mathcal{S}_\tau}\}_{i,j} = \begin{cases} A_{i,j} & \text{if } (i,j) \text{ is an edge of } S_\tau \\ 0 & \text{otherwise.} \end{cases} \tag{3}$$

The thresholded polynomial $P_\tau(X)$ of $P(X)$ with respect to $\mathcal{S}_\tau$ is given by

$$P_\tau(X) = \sum_{n=0}^{M} c_n T^{(n)}_{\mathcal{S}_\tau}(X), \tag{4}$$

where the thresholded $T^{(n)}_{\mathcal{S}_\tau}(X)$ can be calculated from a linear recurrence

$$T^{(n)}_{\mathcal{S}_\tau}(X) = \alpha_n \lfloor X T^{(n-1)}_{\mathcal{S}_\tau}(X) \rfloor_{\mathcal{S}_\tau} + \sum_{m=0}^{n-1} \alpha_m T^{(m)}_{\mathcal{S}_\tau}(X), \tag{5}$$

with $T^{(0)}_{\mathcal{S}_\tau}(X) = I$. A key observation of this paper is that the calculation of $P_\tau(X)$ in Eqs. (4) and (5) is equivalent to a partitioned subgraph expansion on $\mathcal{S}_\tau$. This approach is

illustrated in Fig. 1. For any vertex i of $\mathcal{S}_\tau$, let s_τ^i be the subgraph of $\mathcal{S}_\tau$ induced by the *core* (meaning belonging to a single subgraph) vertex i and all *halo* (shared) vertices that are directly connected to i in $\mathcal{S}_\tau$. Then the ith matrix column of $P_\tau(X)$ is given by the thresholded expansion determined by s_τ^i only, i.e.

$$\{P_\tau(X)\}_{:,i} = \left\{P(\mathbf{x}[s_\tau^i])\right\}_{:,j} \ . \tag{6}$$

Here j is the column (or row) of the polynomial for the subgraph s_τ^i containing all edges from the core vertex i that corresponds to column i of the complete matrix polynomial on the left-hand side. $\mathbf{x}[s_\tau^i]$ is the small dense principal submatrix that contains only the entries of X corresponding to s_τ^i. The full matrix $P_\tau(X)$ can then be assembled, column by column, from the set of smaller dense matrix polynomials $P(\mathbf{x}[s_\tau^i])$ for each vertex i. The calculation of a numerically thresholded matrix polynomial $P_\tau(X)$ thus can be replaced by a sequence of fully independent small dense matrix polynomial expansions determined by a graph partitioning.

Equation (6) represents an exact relation between a globally thresholded sparse matrix algebra and a graph partitioning approach, which is valid for a general matrix polynomial $P(X)$. Several observations can be made about this equivalence: *i*) $P_\tau(X)$ is not symmetric and with the order of the matrix product for the threshold in Eq. (5) we collect $P_\tau(X)$ column by column in Eq. (6) as illustrated by the *directed* graph at the bottom of Fig. 1; *ii*) the accuracy of the matrix polynomial increases (decreases) as the threshold τ is reduced (increased) and the number of edges of $\mathcal{S}_\tau$ increases (decreases); *iii*) we may thus include additional edges in $\mathcal{S}_\tau$ without loss of accuracy; *iv*) the polynomial $P_\tau(X)$ is zero at all entries outside of $\mathcal{S}_\tau$; *v*) apart from spurious cancellations, the non-zero pattern of $P_\tau(X)$ is therefore the same as $\mathcal{S}_\tau$ and we can expect a numerically thresholded exact matrix polynomial, $\lfloor P(X) \rfloor_\tau$, to have a non-zero structure similar to $\mathcal{S}_\tau$; *vi*) the graph partitioning can be generalized such that each vertex corresponds to a combined set of vertices, i.e. a community, without loss of accuracy; *vii*) we may reduce the computational cost by identifying such communities using graph partitioning schemes, which can be tailored for optimal platform-dependent performance; *viii*) the exact relation given by Eqs. (4) - (6) holds for any structure of $\mathcal{S}_\tau$ and is not limited to the threshold in Eq. (2); *ix*) the particular sequence of matrix operations in the calculation of $P_\tau(X)$ is of importance because of the thresholding in Eq. (5), whereas the order (or grouping) of the matrix multiplications is arbitrary for the contracted

matrix polynomials $P(\mathbf{x}[s_\tau^i])$ in Eq. (6); and x) the computational cost of each polynomial expansion is dominated by separate sequences of dense matrix-matrix multiplication that can be performed independently and in parallel.

A main point of this paper is that the equivalence between the calculation of the thresholded sparse matrix polynomial and the graph partitioned expansion in Eq. (6) provides a natural framework for a graph-based formulation of linear scaling electronic structure theory. In Kohn-Sham density functional theory the matrix polynomial in Eq. (1) is replaced by the Fermi-operator expansion[3,51,52] where

$$P(H) = D = \left[e^{\beta(H-\mu)} + 1\right]^{-1} \approx \sum_{n=0}^{M} c_n T^{(n)}(H). \tag{7}$$

Here D is the density matrix, H the Hamiltonian, μ the chemical potential, and β the inverse temperature. The matrix functions, $T^{(n)}(X)$, are typically Chebyshev polynomials constructed by a recurrence equation as in Eq. (5). With a local basis set H and $P(H)$ have sparse matrix representations above some numerical threshold for sufficiently large non-metallic systems[5,6]. The graph-based construction of sparse matrix polynomials in Eq. (6) can then be applied to the calculation of the density matrix with the data dependency graph $\mathcal{S}_\tau$ estimated from an approximate prior density matrix that is available in an iterative self-consistent field (SCF) optimization or from previous time steps in a MD simulation. The computation can be accelerated with a recursive Fermi-operator expansion[32–37,39–41]. In the zero temperature limit the Fermi function equals the Heaviside step function θ and a recursive expansion is then given by $D = \theta(\mu I - H) = \lim_{n\to\infty} f_n(f_{n-1}(\dots f_0(H)\dots))$, which reaches a high expansion order much more rapidly compared to the serial form in Eq. (1). With $f_n(X)$ being 2nd-order polynomials[35] we reach an expansion order of over a billion in only 30 iterations. The ability to use a fast recursive expansion is motivated from ix) above, and since any recursive expansion also can be written in the general form of Eq. (1). Once the density matrix D is known, the expectation value of any operator A is given by $\langle A \rangle = Tr[DA]$. Generalizations to quantum perturbation theory are straightforward[53,54].

B. Macromolecular test system

Figure 2 shows the error per atom in the density matrix of the band energy, $E_{\text{band}} = Tr[DH]$, calculated with the graph-based formulation above for a 19,945-atom macromolec-

ular system of polyalanine solvated in water. The calculations were performed using self-consistent charge density functional tight-binding theory[55–57] as implemented in the electronic structure program LATTE[58] in combination with the recursive second-order spectral projection (SP2) zero-temperature Fermi-operator expansion scheme[35]. The data dependency graphs, $\mathcal{S}_\tau$, were estimated by thresholding an "exact" density matrix with varying thresholds, τ. Different numbers of subgraph communities (512, 1024 or 2048) were chosen and optimized with the METIS heuristic multilevel graph partitioning package[59] for the different data dependency graphs (one for each threshold) using the multilevel recursive bisection method. The errors were determined in comparison to the "exact" density matrix, which was calculated using regular sparse matrix algebra with a tight threshold of 10^{-12}. The error is fairly insensitive to the number of graph partitions and is instead controlled by the value of the threshold that is used to estimate the data dependency graphs. In contrast, the computational cost varies significantly with the size of the graph partitions. The cost in the limit of only one large community, containing the whole system, or in the opposite limit, with one partition for each orbital, scale as $\mathcal{O}(N^3)$ or $\mathcal{O}(Nm^3)$, respectively, where m is the average number of edges per vertex in $\mathcal{S}_\tau$ and $N \times N$ is the size of H. A straightforward graph partitioning may thus lead to a significant overhead compared to a Fermi-operator expansion using thresholded sparse matrix algebra[5], which scales as $\mathcal{O}(Nm^2)$. However, with an optimized graph partitioning the total cost can be reduced to scale as $\mathcal{O}(Nm^2)$. A similar optimization can be performed for divide and conquer methods, but may not be applicable to inhomogeneous systems[17]. Figure 3 shows the timing (66 s, red dashed line) for a thresholded sparse matrix algebra (SpM Alg) Fermi-operator expansion with Intel's MKL sparse matrix library[30] running in parallel on a dual eight-core CPU. With the graph-based approach (filled circles) using the METIS graph partitioning (Graph Part.) program for varying numbers of communities it is possible to significantly reduce the run time on the same platform (23 s). The graph-based formalism also has the additional advantage of an almost trivial and highly scalable parallelism as is demonstrated by the run times on 1, 16 or 32 graphics processing units (GPUs) on separate nodes (open symbols)[60]. The parallel performance is close to ideal, reaching a performance of about 25 μs/atom and a subsecond wall-clock time on the 32 node GPU platform.

C. Molecular dynamics simulation

Linear scaling divide and conquer methods[12–17] rely on an estimated finite range of direct electron interaction, which can be motivated by the localized character of the Wannier functions[61–63]. This allows a system to be partitioned into smaller overlapping regions that are solved separately (apart from long-range electrostatic interactions), within predetermined local interaction zones, and then reassembled. Divide and conquer schemes are naturally parallel and in spirit similar to our graph-based approach. However, their numerical accuracy can be difficult to control without careful prior testing and convergence analysis[6,64,65]. An automatic, adjustable error control is particularly challenging in MD simulations of inhomogeneous materials, where reacting or floppy molecules and atoms can move across pre-determined local interaction zones and where transitions between localized and itinerant electronic states may occur. MD simulations of inhomogeneous molecular systems with significant changes in the electronic overlap are therefore of particular interest when we evaluate our framework. Furthermore, the precision can be gauged very sensitively by the accuracy and long-term stability of the total energy.

The data dependency graph $\mathcal{S}_\tau(t)$ can be estimated from the numerically thresholded density matrix in the previous MD time step, $\lfloor D(t-\delta t)\rfloor_\tau$, and new Hamiltonian matrix elements, $H(t)$, as the atoms move, for example, from

$$\mathcal{S}_\tau(t) \leftarrow \lfloor(\lfloor D(t-\delta t)\rfloor_\tau + H(t))^2\rfloor_\epsilon. \tag{8}$$

In our MD simulation below we use the symbolic representation of $\mathcal{S}_\tau(t)$ in Eq. (8), which is given from the non-zero pattern of the thresholded density matrix (with $\tau = 10^{-4}$) combined with the non-zero pattern of $H(t)$, and instead of the matrix square we use paths of length two, corresponding to the symbolic operation ($\epsilon = 0$). This approach that adapts $\mathcal{S}_\tau(t)$ to each new MD time step by including additional redundant edges works surprisingly well, though $\mathcal{S}_\tau(t)$ cannot increase by more than paths of length two between two MD steps. Generalized estimates can also be designed for the iterative SCF optimization.

Figure 4 shows the fluctuations of the total energy during a microcanonical MD simulation of liquid water that was performed using LATTE[58] and the extended Lagrangian formulation of Born-Oppenheimer MD[50,66–69]. The density matrix was calculated from a partitioning over separate subgraphs of $\mathcal{S}_\tau(t)$, with one water molecule per core. For the Fermi-operator expansion (at zero temperature) we used the recursive SP2 algorithm[35]. In

each time step the complete SP2 sequence (the same for each subgraph expansion) for the correct total occupation is pre-determined from the HOMO-LUMO gap that is estimated from the previous time step as in Ref.[41]. In this way each expansion can be performed independently, without exchange of information during or between the matrix multiplications as otherwise would be required[8,28]. Communication is reduced to a minimum and no additional adjustments of the electronic occupation, as in divide and conquer calculations[14], is required. The inset of Fig. (4) shows the number of water molecules of a single subgraph (core + halo) along the trajectory of an individual molecule, which oscillates as $\mathcal{S}_\tau(t)$ adaptively follows the fluctuations in the electronic overlap. Despite the large oscillations, including between 1 and 25 molecules, the total energy is both accurate and stable. The "exact" calculation with fully converged density matrices ($\geq$4 SCFs per step) using dense matrix algebra based on full $\mathcal{O}(N^3)$ diagonalization, is virtually indistinguishable for the first 0.5 ps (or 1,000 time steps).

III. DISCUSSION

Linear scaling QMD simulations using divide and conquer or radial truncation approaches often show systematic energy drifts[70–72] that are significantly higher than regular $\mathcal{O}(N^3)$ methods[9,42,43] and multiple orders of magnitude larger than the graph-based QMD simulation in Fig. 4. Such problems may occur because of difficulties controlling the error in the force evaluations[6,73] as atoms move across the local zone boundaries and as the electronic overlap fluctuates, or because of incomplete SCF optimization causing a broken time-reversal symmetry[42,74]. The problem is illustrated in Fig. 5, which shows a comparison between a divide and conquer approach and our graph-based calculation of the density matrix for a snapshot from a MD simulation of the water system in Fig. 4. Without the adaptivity of the graph-based method, the divide and conquer approach needs a large cutoff radius, $R_{\rm cut}$, to reach sufficient convergence in the calculation of the density matrix for the water system, which leads to a significant overhead. With the graph-based framework as demonstrated here in combination with a modern formulation of Born-Oppenheimer MD[42–50], these problems can be avoided.

The off-the-shelf METIS graph partitioning scheme used for the macromolecular system in Figs. 2 and 3 works very well and drastically reduces the overhead compared to

a straightforward implementation. However, by adjusting the graph partitioning to the particular requirements of the electronic structure calculation as well as the computational platform, further optimizations are possible[75]. We may also incorporate an on-the-fly graph partitioning that updates the subgraphs during a MD simulation as the structure changes.

IV. CONCLUSIONS

In this article we have shown how graph theory can be combined with quantum theory to calculate the electronic structure of large complex systems with well-controlled accuracy. The graph formalism is general and applicable to a broad range of electronic structure methods and materials, for which sparse matrix representations can be used, including QMD simulations, overcoming significant gaps in linear scaling electronic structure theory.

ACKNOWLEDGMENTS

We acknowledge support from Office of Basic Energy Sciences (LANL2014E8AN) and the Laboratory Directed Research and Development program of Los Alamos National Laboratory (LANL). Generous support and discussions with T. Peery at the T-division International Java Group are acknowledged. The research used resources provided by the LANL Institutional Computing Program. LANL, an affirmative action/equal opportunity employer, is operated by Los Alamos National Security, LLC, for the National Nuclear Security Administration of the U.S. DOE under Contract No. DE-AC52-06NA25396.

---

[1] P. Hohenberg and W. Kohn, Phys. Rev., **136**, B:864 (1964).

[2] W. Kohn and L. J. Sham, Phys. Rev. B, **140**, A1133 (1965).

[3] R. G. Parr and W. Yang, *Density-functional theory of atoms and molecules* (Oxford University Press, Oxford, 1989).

[4] R. Dreizler and K. Gross, *Density-functional theory* (Springer Verlag, Berlin Heidelberg, 1990).

[5] S. Goedecker, Rev. Mod. Phys., **71**, 1085 (1999).

[6] D. R. Bowler and T. Miyazaki, Rep. Prog. Phys., **75**, 036503 (2012).

[7] D. R. Bowler and T. Miyazaki, J. Phys.: Condens. Matter, **22**, 074207 (2010).

[8] J. VandeVondele, U. Borstnik, and J. Hutter, J. Chem. Theory Comput., **8**, 3565 (2012).

[9] D. Marx and J. Hutter, "Modern methods and algorithms of quantum chemistry," (ed. J. Grotendorst, John von Neumann Institute for Computing, Jülich, Germany, 2000) 2nd ed.

[10] G. Chartrand, *Introductory Graph Theory* (Dover Publications, New York, 1985).

[11] J. A. Bondy, *Graph Theory* (Springer-Verlag, London, 2008).

[12] W. Yang, Phys. Rev. Lett., **66**, 1438 (1991).

[13] P. D. Walker and P. G. Mezey, J. Am. Chem. Soc., **115**, 12423 (1993).

[14] W. T. Yang and T. S. Lee, J. Chem. Phys., **103**, 5674 (1995).

[15] I. A. Abrikosov, A. M. N. Niklasson, S. I. Simak, B. Johansson, A. V. Ruban, and H. L. Skriver, Phys. Rev. Lett., **76**, 4203 (1996).

[16] K. Kitaura, E. Ikeo, T. Nakano, and M. Uebayasi, Chem. Phys. Lett., **313**, 701 (1999).

[17] T. Ozaki, Phys. Rev. B, **74**, 245101 (2006).

[18] F. G. Gustavson, ACM Trans. Math. Soft., **4**, 250 (1978).

[19] S. Pissanetzky, *Sparse Matrix Technology* (Academic Press, London, 1984).

[20] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipies in FORTRAN* (Cambridge University Press, Port Chester, NY, 1992).

[21] Y. Saad, *Iterative methods for Sparse Linear Systems* (PWS Publishing, Boston, 1996).

[22] M. Challacombe, Comput. Phys. Commun., **128**, 93 (2000).

[23] E. H. Rubensson, E. Rudberg, and P. Salek, J. Comput. Chem., **28**, 2531 (2007).

[24] E. H. Rubensson, E. Rudberg, and P. Salek, J. Chem. Phys., **128**, 74109 (2008).

[25] A. Buluc and J. R. Gilbert, SIAM J. Sci. Comput., **34**, 170 (2012).

[26] U. Borstnik, J. VandeVondele, V. Weber, and J. Hutter, Parallel Computing, **40**, 47 (2014).

[27] N. Bock, M. Challacombe, and L. V. Kale, arXiv http://arxiv.org/abs/1403.7458v5 (2013).

[28] V. Weber, T. Latino, A. Pozdeev, I. Feduova, and A. Curioni, J. Chem. Theory Comput., **11**, 3145 (2015).

[29] S. M. Mniszewski, M. J. Cawkwell, M. E. Wall, J. Mohd-Yusof, N. Bock, T. C. Germann, and A. M. N. Niklasson, J. Chem. Theory Comput., **11**, 4644 (2015).

[30] Intel MKL, Intel Math Kernel Library, https://software.intel.com/en-us/intel-mkl (2015).

[31] NVIDIA cuSPARSE, https://developer.nvidia.com/cusparse (2014).

[32] R. McWeeny, Proc. R. Soc. London Ser. A-Math, **235**, 496 (1956).

[33] A. H. R. Palser and D. E. Manolopoulos, Phys. Rev. B, **58**, 12704 (1998).

[34] A. Holas, Chem. Phys. Lett., **340**, 552 (2001).

[35] A. M. N. Niklasson, Phys. Rev. B, **66**, 155115 (2002).

[36] A. M. N. Niklasson, Phys. Rev. B, **68**, 233104 (2003).

[37] W. Z. Liang, C. Saravanan, Y. Shao, R. Baer, A. T. Bell, and M. Head-Gordon, J. Chem. Phys., **119**, 4117 (2003).

[38] E. Rudberg and E. H. Rubensson, J. Phys.: Condens. Matter, **23**, 075502 (2011).

[39] E. H. Rubensson, J. Chem. Theory and Comput., **7**, 1233 (2011).

[40] P. Suryanarayana, Chem. Phys. Lett., **555**, 291 (2013).

[41] E. H. Rubensson and A. M. N. Niklasson, SIAM J. Sci. Comput., **36**, 148 (2014).

[42] P. Pulay and G. Fogarasi, Chem. Phys. Lett., **386**, 272 (2004).

[43] J. Herbert and M. Head-Gordon, Phys. Chem. Chem. Phys., **7**, 3269 (2005).

[44] A. M. N. Niklasson, C. J. Tymczak, and M. Challacombe, Phys. Rev. Lett., **97**, 123001 (2006).

[45] T. D. Kühne, M. Krack, F. R. Mohamed, and M. Parrinello, Phys. Rev. Lett., **98**, 066401 (2007).

[46] G. Zheng, A. M. N. Niklasson, and M. Karplus, J. Chem. Phys., **135**, 044122 (2011).

[47] J. Hutter, WIREs Comput. Mol. Sci., **2**, 604 (2012).

[48] L. Lin, J. Lu, and S. Shao, Entropy, **16**, 110 (2014).

[49] M. Arita, D. R. Bowler, and T. Miyazaki, J. Chem. Theory Comput., **10**, 5419 (2014).

[50] A. M. N. Niklasson and M. Cawkwell, J. Chem. Phys., **141**, 164123 (2014).

[51] S. Goedecker and L. Colombo, Phys. Rev. Lett., **73**, 122 (1994).

[52] R. N. Silver and H. Roder, Int. J. Mod. Phys. C, **5**, 735 (1994).

[53] A. M. N. Niklasson and M. Challacombe, Phys. Rev. Lett., **92**, 193001 (2004).

[54] V. Weber, A. M. N. Niklasson, and M. Challacombe, Phys. Rev. Lett., **92**, 193002 (2004).

[55] M. Elstner, D. Poresag, G. Jungnickel, J. Elsner, M. Haugk, T. Frauenheim, S. Suhai, and G. Seifert, Phys. Rev. B, **58**, 7260 (1998).

[56] M. W. Finnis, A. T. Paxton, M. Methfessel, and M. van Schilfgarde, Phys. Rev. Lett., **81**, 5149 (1998).

[57] T. Frauenheim, G. Seifert, M. E. aand Z. Hajnal, G. Jungnickel, D. Poresag, S. Suhai, and R. Scholz, Phys. Stat. sol., **217**, 41 (2000).

[58] M. J. Cawkwell and A. M. N. Niklasson, J. Chem. Phys., **137**, 134105 (2012).

[59] G. Karypis, , and V. Kumar, SIAM J. Sci. Comput., **20**, 359 (1999).

[60] NVIDIA cuBLAS, https://developer.nvidia.com/cuBLAS (2014).

[61] W. Kohn, Phys. Rev. Lett., **76**, 3168 (1996).

[62] W. Kohn, Phys. Rev. A, **133**, A171 (1964).

[63] N. F. Mott, Phil. Mag., **6**, 278 (1961).

[64] T. S. Lee, D. M. York, and W. Yang, J. Chem. Phys., **105**, 2744 (1996).

[65] D. M. York, T. S. Lee, and W. Yang, Phys. Rev. Lett., **80**, 5011 (1998).

[66] A. M. N. Niklasson, Phys. Rev. Lett., **100**, 123004 (2008).

[67] P. Steneteg, I. A. Abrikosov, V. Weber, and A. M. N. Niklasson, Phys. Rev. B, **82**, 075110 (2010).

[68] P. Souvatzis and A. M. N. Niklasson, J. Chem. Phys., **140**, 044117 (2014).

[69] B. Aradi, A. M. N. Niklasson, and T. Frauenheim, J. Chem. Theory Comput., **11**, 3357 (2015).

[70] F. Shimojo, R. K. Kalia, A. Nakano, and P. Vashista, Phys. Rev. B, **77**, 085103 (2008).

[71] E. Tsuchida, J. Phys.: Condens. Matter, **20**, 294212 (2008).

[72] F. Shimojo, S. Hattori, R. K. Kalia, M. Kusaneth, W. W. Mou, A. Nakano, K. Nomura, S. Ohmura, P. Rajak, K. Shimamura, and P. Vashista, J. Chem. Phys., **140**, 18529 (2014).

[73] M. Kobayashi, T. Kunisada, T. Akama, D. Sakura, and H. Nakai, J. Chem. Phys., **134**, 034105 (2011).

[74] D. K. Remler and P. A. Madden, Mol. Phys., **70**, 921 (1990).

[75] H. Djidjev and et al., (2016), unpublished.

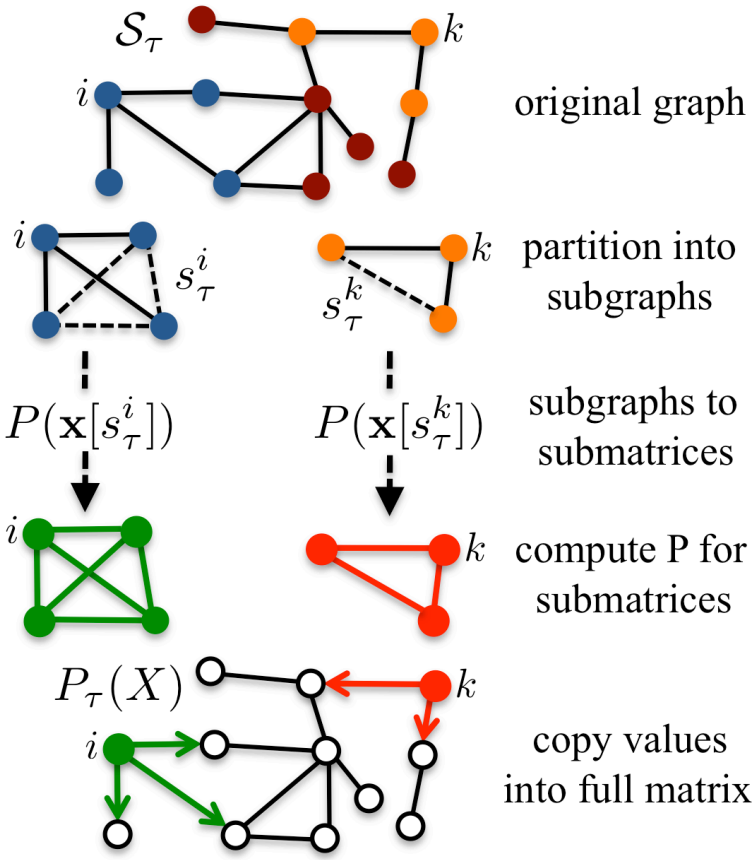

FIG. 1. The data dependency graph $\mathcal{S}_\tau$ and the subgraphs (s^i_τ or s^k_τ), one for each *core* vertex (i or k) including all directly connected *halo* vertices in $\mathcal{S}_\tau$. The full matrix polynomial $P_\tau(X)$ is given by an assembly from $P(\mathbf{x}[s^i_\tau])$ of the separate dense subgraph contractions $\mathbf{x}[s^i_\tau]$.

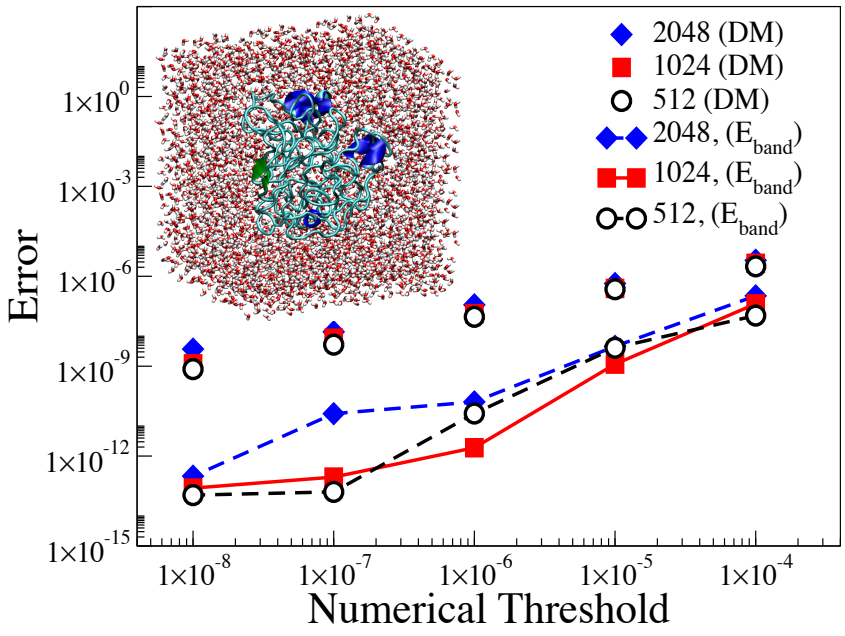

FIG. 2. The error in the calculated density matrix (DM) for polyalanine (2,593 atoms) in water with a total of 19,945 atoms (inset), as measured by the Frobenius norm (normalized per atom) for partitions with 512, 1024, and 2048 separate communities based on graphs, $\mathcal{S}_\tau$, from varying numerical thresholds τ. The connected symbols (lower part) shows the error in band energy, $E_{\text{band}} = Tr[HD]$, per atom.

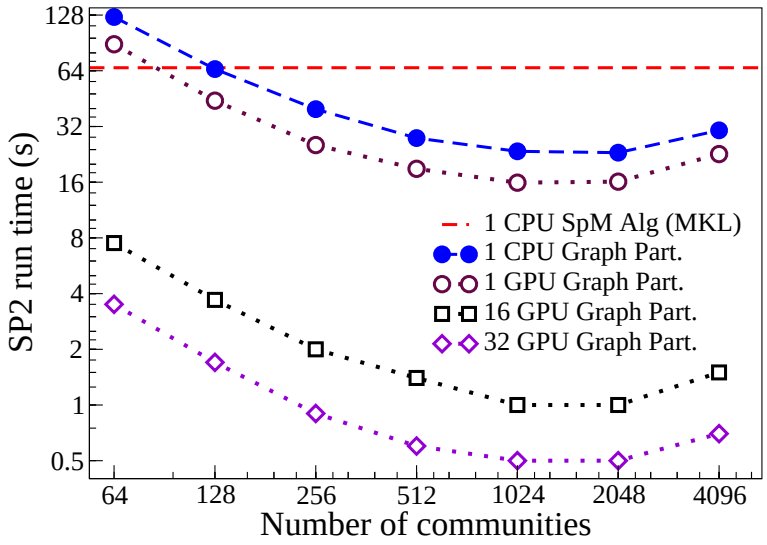

FIG. 3. The time to calculate the density matrix using the SP2 expansion (with threshold $\tau = 10^{-5}$) partitioned over different sets of subgraphs for the solvated polyalanine system (19,945 atoms).

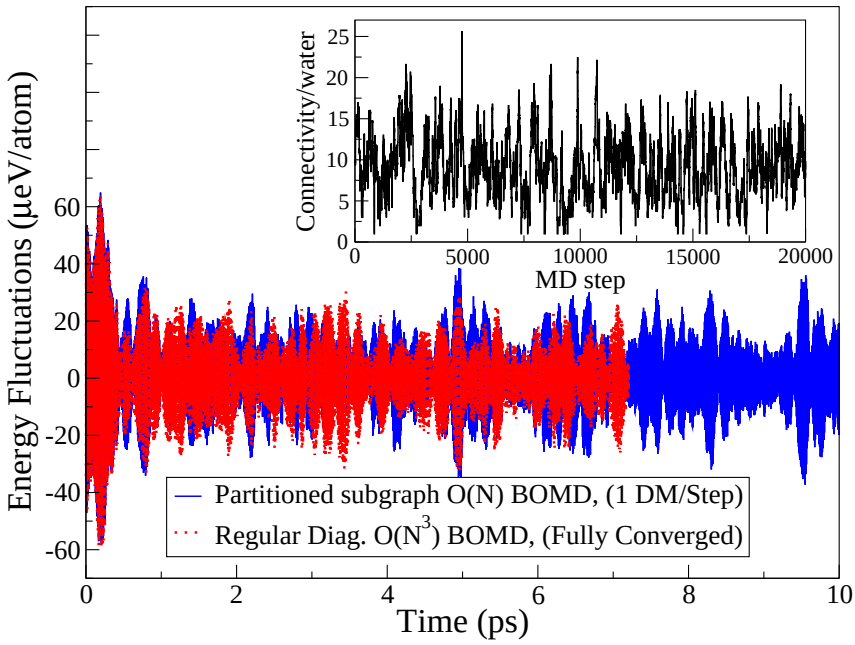

FIG. 4. The total energy fluctuations in a microcanonical Born-Oppenheimer MD (BOMD) simulation of liquid water (100 molecules, $T \sim 300$ K, $\delta t = 0.5$ fs), using graph partitioning and one density matrix (DM) construction per step vs. SCF optimized BOMD with diagonalization (Diag.). The inset shows the number of water molecules associated with the subgraph of an individual molecule. Energy drift is less than $\sim 0.2\mu$eV/atom per ps.

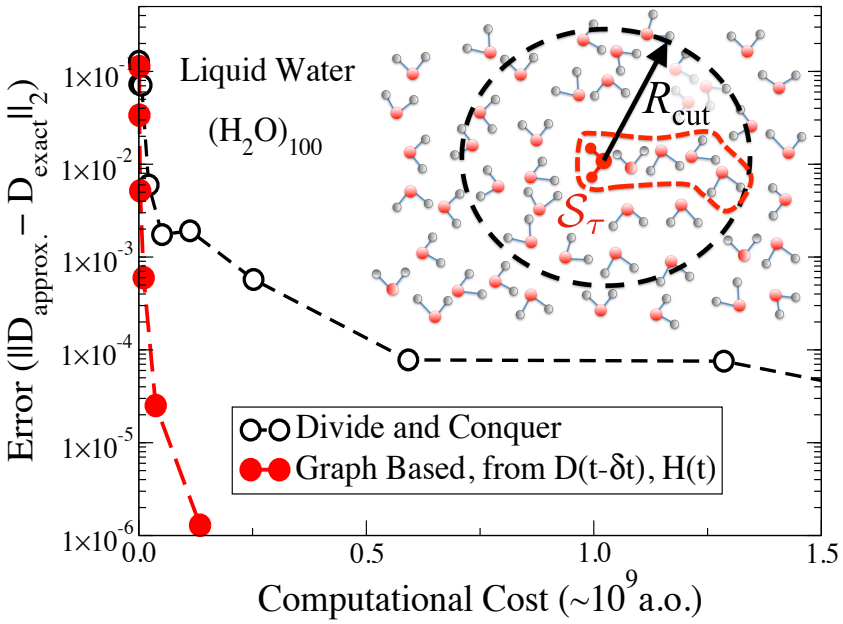

FIG. 5. The convergence of the density matrix error for a snapshot during a MD simulation of the water system in Fig. 4 (100 molecules, $T \sim 300$ K, $\delta t = 0.5$ fs) as a function of the computational cost for various numerical thresholds ($\tau = 10^{-1}, 10^{-2}, \ldots, 10^{-6}$) in the symbolic estimate of the data-dependency graph in Eq. (8) for the graph-based method, and for different sizes of the cutoff radius, R_{cut}, in a divide and conquer approach. To capture a hypothetical electronic overlap within the red dashed border in the inset (associated with the data-dependency graph $\mathcal{S}_\tau$ for the large red molecule at the center), the cutoff radius needs to be large, which leads to a significant overhead for the divide and conquer approach. The efficiency would be similar only for a homogeneous system. The computational cost was estimated from the sum of the number arithmetic operations (a.o.) required to calculate the density matrices from all the separate subgraph partitions or divide and conquer regions – one for each water molecule.